

Object detection and amplitude estimation based on maximum *a posteriori* reconstructions

K. M. Hanson*

Los Alamos National Laboratory, MS P940
Los Alamos, NM 87545 USA

Abstract

We report on the behavior of the linear maximum *a posteriori* (MAP) tomographic reconstruction technique as a function of the assumed rms noise σ_n in the measurements, which specifies the degree of confidence in the measurement data. The unconstrained MAP reconstructions are evaluated on the basis of the performance of two related tasks; object detection and amplitude estimation. It is found that the detectability of medium-sized discs remains constant up to relatively large σ_n before slowly diminishing. However, the amplitudes of the discs estimated from the MAP reconstructions increasingly deviate from their actual values as σ_n increases.

Introduction

We have previously presented a method to test image-recovery algorithms on the basis of how well the resulting reconstructions allow one to perform specific tasks set forth for the imaging system [1,2]. Task performance is numerically evaluated for a specified imaging situation by a Monte Carlo simulation of the entire imaging process including random scene generation, data taking, reconstruction, and task performance. An essential aspect of this method is that the evaluation is based on many randomly generated scenes leading to a statistically significant estimate of performance. In our previous work we have concentrated on the ART algorithm, not for any strong reason, but because it demonstrates good convergence properties particularly in the presence of constraints, such as that of nonnegativity.

Here we look at the behavior of the linear maximum *a posteriori* (MAP) reconstruction method, which is derived assuming that the *a priori* probability distribution for the ensemble of images under consideration can be described by a multivariate normal distribution. The essence of the Bayesian approach is the use of prior information expressed in a probabilistic sense concerning the class of images being reconstructed. The influence of the prior knowledge on the solution is controlled by the assumed rms noise σ_n in the data. The MAP estimate is known to be unbiased, that is, its mean value equals the true value, when averaged over the complete ensemble of images [3]. For each individual scene, however, the MAP reconstruction deviates from the actual image towards \bar{f} . This systematic deviation grows as σ_n increases.

In this paper we address the performance of two visual tasks based on images reconstructed from projections using the linear MAP reconstruction algorithm: object detection and object amplitude estimation. These tasks are related because the best amplitude estimate is the appropriate decision variable for the detection task. We observe different behavior as a function of σ_n , the assumed value for the rms noise in the projection measurements, for the accuracy of the performance of these two tasks based on MAP reconstructions. The choice of σ_n affects the declination in the estimated amplitude from its proper value and the amount of blurring of the reconstructed image. Such a choice is required in a wide variety of algorithms to control the ill conditioning of the inversion process, as in the stopping rule required in some iterative algorithms such as the estimation-maximization (EM) algorithm [4], the algebraic reconstruction technique (ART) [5], and the simultaneous reconstruction technique (SIRT) [6]. Equivalently, χ^2 , which is the sum of the squares of each residual divided by the assumed rms noise value for that measurement, can

*This work was supported by the United States Department of Energy under contract number W-7405-ENG-36.

be used as a side constraint. Such a constraint is employed in the constrained least-squares algorithm of Hunt [7] and the maximum entropy (MAXENT) reconstruction algorithm [8]. A prevalent rule-of-thumb is that χ^2 should equal the number of measurements [7,4]. This constraint might be unsettling to those who learned that in parameter estimation, the best estimate is taken to be at the minimum in χ^2 [9]. As χ^2 has its foundation in the assumed normal probability distribution for the measurement errors, the minimum χ^2 condition is the same as that of maximum likelihood. It is known that *if* there exists an unbiased solution with minimum variance, it is the same as the maximum-likelihood solution [10]. Thus there seems to be some justification for proceeding to the minimum χ^2 and not stopping at a predetermined value. How then do we choose the best value of χ^2 (or σ_n) in MAP? The numerical evaluation procedure employed here provides one method to investigate how this choice affects task performance based on the MAP reconstructions.

Linear MAP Reconstruction

The linear maximum *a posteriori* (MAP) algorithm is derived under the assumption that the *a priori* probability distribution for the ensemble of images under consideration is described by a multivariate normal distribution. The relevant parameters used in that description are the vector \bar{f} , the mean of all images in the ensemble, and R_f , the ensemble covariance matrix. The Bayesian approach then leads to the requirement that the solution minimizes [11]

$$\phi(f) = (f - \bar{f})^T R_f^{-1} (f - \bar{f}) + (g - Hf)^T R_n^{-1} (g - Hf) , \quad (1)$$

where g is the measurement vector, related to the image vector f by the measurement equation

$$g = Hf + n , \quad (2)$$

where n is the additive measurement noise vector and R_n the noise covariance matrix. In computed tomography (CT), the H matrix represents the complete set of available projection measurements and matrix H^T the backprojection process. The second term of (1) is derived from the expression for the likelihood, which is based on the assumption that the measurement noise follows a normal (Gaussian) probability law. The influence of the prior knowledge on the solution is controlled by the ratio of the ensemble variance to the noise variance σ_n^2 . As Eq. (1) is a simple quadratic form, it is guaranteed to have a unique solution, in the absence of nonlinear constraints.

The solution to (1) is found by setting the gradient of ϕ to zero, leading to the linear MAP equation:

$$R_f^{-1} (\bar{f} - \hat{f}_{MAP}) + H^T R_n^{-1} (g - H \hat{f}_{MAP}) = 0 . \quad (3)$$

It is observed that the MAP solution \hat{f}_{MAP} is a balance between \bar{f} and the solution to the measurement equation (2), which is also called the maximum-likelihood solution. The MAP estimate is known to be unbiased, that is, when averaged over the complete ensemble of images, the estimate is the same as the average true function value [3]. This statement is true because the first term of Eq. (3) effectively pulls the estimate towards the ensemble mean \bar{f} . However, this statement is misleading because it only refers to the average value taken over the ensemble of images. If the average were carried out over all noise realizations for a single scene, the MAP reconstruction would deviate from the actual image towards \bar{f} . This deviation grows as σ_n increases. As a fundamental tenet, we maintain that what is really important is how well tasks can be performed on basis of the resulting reconstructions. The abovementioned effects may or may not be important in the performance of relevant visual tasks. Hence, it is important to directly evaluate how well such tasks may be performed.

For stationary blur matrices (H) and covariance matrices, the usual approximation of replacing Toeplitz matrices by circulants [7] transforms Eq. (3) into the following expression in the Fourier domain:

$$\hat{f}_{MAP} = \left[\frac{\eta}{H^*H + \eta} \right] \bar{f} + \left[\frac{H^*}{H^*H + \eta} \right] g , \quad (4)$$

where $\eta = \sigma_n^2/\sigma_f^2$, which reflects the assumption that the covariance matrices lack cross correlations so that $R_f = \text{diag}(\sigma_f^2 I)$ and $R_n = \text{diag}(\sigma_n^2 I)$. In this expression, H is the Fourier amplitude corresponding to the blur function and H^* is its complex conjugate. All quantities (except η) are implicit functions of the 2D spatial frequency. Equation (4) is the Wiener filter with the addition of the leading bias term, proportional to \bar{f} , the ensemble average.

The limiting forms of this expression are of interest. For very large σ_n :

$$\hat{f}_{MAP} \approx (1 - \eta^{-1}H^*H)\bar{f} + \eta^{-1}H^*g, \quad (5)$$

which approaches \bar{f} plus a term whose contrast is decreasing and is proportional to a blurred version of difference between the data g and the measurements corresponding to \bar{f} . Ignoring the noise, the last term is proportional to H^*Hf , that is, a doubly blurred version of the original image. Therefore, as σ_n grows, the reconstruction approaches the ensemble average image with a low-contrast, blurred version of the difference of the original image from \bar{f} added to it. The blurring degrades task performance on two accounts. First, it mixes noise with signal and, second, different objects are blended together, confusing the interpretation of any individual object. We observe from (3) or (4) that between the two limits the linear MAP solution responds in a continuous fashion to η .

As $\sigma_n \rightarrow 0$:

$$\lim_{\eta \rightarrow 0} \hat{f}_{MAP} = \delta(|H|)\bar{f} + \lim_{\eta \rightarrow 0} \left[\frac{H^*}{H^*H + \eta} \right] g. \quad (6)$$

Note that at the frequencies where $|H| = 0$, the second term is zero and $\hat{f}_{MAP} = \bar{f}$. Thus, where the measurements provide no information, the estimate takes on the default value \bar{f} . In terms of the space of reconstructed functions, this region is called the null space [11]. The second term is recognized as the Fourier transform of $H^\dagger g$, where H^\dagger is the Moore-Penrose pseudoinverse of the matrix H [12,7]. This term is the same as the familiar expansion of the pseudoinverse of the matrix H in terms of singular value decomposition in which the terms with zero singular value are dropped. When $\bar{f} = 0$ where $|H| = 0$, Eq. (6) is the maximum-likelihood estimate, which is also the least-squares or minimum χ^2 solution. This type of solution is known to often be ill conditioned, that is, unstable and noisy because the matrix $H^T H$ is singular or nearly so. The various techniques for ameliorating solutions to ill-posed problems are brought into a general framework by regularization theory [13,14]. In that context the MAP approach overcomes ill conditioning through the regularization provided by the first term of Eq. (3).

In the generalized Tikhonov approach to regularization [13], the solution vector f is given by the minimum of

$$|g - Hf|^2 + \lambda\Omega(f), \quad (7)$$

where $\lambda \geq 0$ and $\Omega(f)$ is a positive scalar function of f . The solution strikes a balance between a least-squares solution (maximum-likelihood estimate) to the measurement equations $g = Hf$ and the minimizer of the regularization term $\Omega(f)$, whose strength is controlled by the scalar λ . The regularization term is included to transform the often ill-posed problem of finding the least-squares solution into one that is well posed. At the same time it allows one to impose desired properties on the solution, which may be (but often are not) founded on prior knowledge about the class of f s under consideration.

The similarity between (7) and (1) suggests a common behavior of the MAP method and the method of Tikhonov regularization. Thus the effect of the use of prior information, which is the hallmark of the MAP approach, is to provide regularization. The stability of the MAP solution may be analyzed with the tools provided by regularization theory [14,15]. Conversely, Tikhonov regularization may be interpreted in terms of the Bayesian approach. The regularization term can be chosen or even judged on the basis of what is reasonable for the kinds of images with which one is dealing in any specific application. A mathematically sound basis for choosing the regularization term might be found in the ensemble properties of the images in question. However, the best performance of specified tasks requires a more complete analysis of the properties of the ensemble than the simple mean bias.

Iterative Solution Algorithm

An iterative method provides a convenient way of finding the solution to the MAP equation [11]. The initial estimate is the ensemble mean image

$$\hat{f}^0 = \bar{f}. \quad (8)$$

The residual of the linear MAP equation (3) for the k th estimate is

$$r^k = \bar{f} - \hat{f}^k + R_f H^T R_n^{-1} (g - H \hat{f}^k) \quad (9)$$

$$= R_f H^T R_n^{-1} g + \bar{f} - A \hat{f}^k, \quad (10)$$

where the matrix A is defined as

$$A = (I + R_f H^T R_n^{-1} H). \quad (11)$$

In the previous iterative scheme [16,11], this residual was used to update \hat{f}^k . As r^k is just the gradient of the function to be minimized, i. e. $\nabla\phi$, that algorithm is equivalent to the method of steepest descent, which is known to suffer a slow rate of convergence as the minimum is approached. An alternative scheme that converges more quickly is the method of conjugate gradients [17], in which an intermediate vector is formed from r^k :

$$p^k = r^k + \beta^{k-1} p^{k-1}, \quad (12)$$

where by definition, $\beta^0 = \beta^1 = 0$. Scalar β^{k-1} is chosen as

$$\beta^{k-1} = -\frac{(Ar^k, p^{k-1})}{(Ap^{k-1}, p^{k-1})}, \quad (13)$$

where the inner product of vectors a and b is denoted by (a, b) , to make each p^k A -orthogonal to the previous one, i. e. $(Ap^k, p^{k-1}) = 0$. Then the update to \hat{f}^k

$$\hat{f}^{k+1} = \hat{f}^k + \alpha^k p^k, \quad (14)$$

where the scalar

$$\alpha^k = \frac{(r^k, r^k)}{(Ap^k, r^k)}, \quad (15)$$

is chosen to make r^{k+1} orthogonal to r^k . Note that in a single updating procedure corresponding to one complete pass through the data, called one iteration, this algorithm requires six projection (H) or backprojection (H^T) operations. These operations, which take the same amount of time to complete, dominate the calculation time for the reconstruction. The algebraic reconstruction technique (ART) [5] requires only two such operations per iteration, and is therefore three times faster per iteration. While this conjugate-gradient algorithm is found to converge more quickly in unconstrained problems than the preceding one, both algorithms tend to stagnate when the nonnegativity constraint is invoked. Thus no results are presented here for constrained MAP reconstruction.

In our reconstruction procedure, each input projection is routinely slightly blurred before being passed onto the reconstruction algorithm. The blur function is triangular in shape and has a FWHM of 3 projection samples. The purpose of this preblurring process is to make sure the input projections do not contain higher spatial frequencies than can be reproduced by the combination of the pixel representation of the reconstruction and the projection algorithm. Otherwise the reconstruction algorithm tends to induce ringing (overshoot and undershoot in response to rapid transitions) in the reconstruction in an attempt to match the input data. This phenomenon is identical to that of Gibb's [18] as both arise from an insufficient ability of the representation to reproduce the input function. The consequence of the mild preblurring is that the rms value of the noise in the data presented to the reconstruction algorithm is reduced by a calculable factor of $\sqrt{19/9}$ or 0.4843, and the noise becomes slightly correlated. We will quote the rms values of the noise after preblurring.

Method to Evaluate Task Performance

The overall method for evaluating a reconstruction algorithm used here has been described before [1,2]. In this method one numerically evaluates a task performance index for a specified imaging situation. This technique consists of a Monte Carlo simulation of the entire imaging process including random scene generation, data taking, reconstruction, and performance of the specified task. The accuracy of the task performance is determined by comparison of the results with the known original scene using an appropriate figure of merit. Repetition of this process for many randomly generated scenes provides a statistically significant estimate of the performance index. When the ability to perform a task is marginal, task performance is inherently stochastic in nature. This is obvious when the data are significantly degraded by random noise. However, it is also often true in measurement geometries in which the data are limited and noiseless. The artifacts produced by the ambiguities arising from the lack of data depend on the scene. As the scenes vary in an uncontrollable manner, so do the artifacts. Interpretation of the image can thus vary from one scene to the next in an apparently random fashion [2].

As the details of the evaluation technique are given elsewhere, they will not be discussed here. The amplitude estimate for each disc is taken to be the average reconstruction value over the area of the disc. This choice does not yield accurate estimates of the amplitude because it does not take into account the blurring effects of the discs, which may be sizable. For the decision variable ψ , which is used to determine whether or not a disc is present, it is assumed that the same average value is appropriate. The detectability index d' is based on the two frequency distributions of the decision variable ψ calculated where the object is known to be present and where none is present [1,2].

Results

In this study we wish to elucidate the behavior of the linear MAP solution as the assumed rms noise value σ_n is varied. The effect of changing σ_n is to alter the relative weights of the two terms in the linear MAP equation (3), thus changing the strength of the prior. In each example, the actual rms noise value remains fixed, leaving the input data unchanged.

We will use for these examples the same specifications for the scenes to be imaged as we used before [1,19]. The first scene in the sequence of ten test images is displayed in Fig. 1. The test images include ten high-contrast discs with amplitude 1.0 and ten low-contrast discs with amplitude 0.1. The discs, which have a diameter of 8 pixels, are randomly placed without overlap in a circle of diameter 128 pixels. The experimental data in this tomographic problem consist of parallel projections, each view containing 128 samples. Each projection is accurately calculated as a strip integral of the analytic expressions for the discs. As such, the projections include high spatial frequency components, necessitating the preblurring operation described above. Gaussian-distributed pseudorandom noise is added to the projections before preblurring to simulate measurement noise.

In the first measurement scenario, 100 views evenly spaced over 180° are assumed and the rms noise in the preblurred projection data is 1.9. This level compares to the peak projection value for a single low-contrast disc of 0.8. In this measurement situation, the data are sufficient in number to provide an artifact-free reconstruction. The reconstructions obtained with 50 iterations of the conjugate-gradient algorithm are shown in Fig. 2. This number of iterations proved sufficient to solve the linear MAP equation (3) to very good accuracy in all cases but $\sigma_n = 0.1$, which did not fully converge. The covariance matrices are assumed to be proportional to the identity matrix. The ensemble characteristics are accurately reproduced by setting \bar{f} to 0.0423 and the diagonal elements of R_f to σ_f^2 , where $\sigma_f = 0.19$. A reconstruction obtained using 100 iterations of the algebraic reconstruction technique (ART) [5] is also shown. The relaxation parameters $\lambda_0 = 0.75$, $r_\lambda = 0.960$, defined in [1], are purposely chosen to produce a result that closely resembles the MAP reconstruction with $\sigma_n = 0.1$. Great latitude in the character of the ART reconstruction exists through the variation of these relaxation parameters. They can, for example, be chosen to optimize the reconstruction for the performance of specified tasks [20,19,21]. As an aside, their values, as well as the

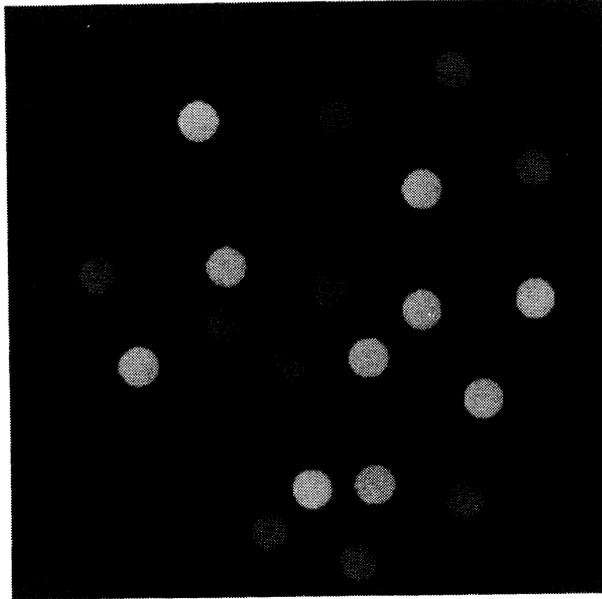


Figure 1: The first in the sequence of ten test images used to evaluate task performance. Two related tasks are to be performed on the low-contrast discs; estimation of their amplitudes and detection of their presence.

number of iterations chosen, relates to the amount of regularization employed in the solution [14], something we will discuss in the last section. Thus, neither of these reconstructions represents the true pseudoinverse solution, but they are approximations to it.

We observe in the reconstructions displayed in Fig. 2 the tendencies discussed earlier. When the assumed rms noise σ_n is very small, the MAP algorithm yields a reconstruction very similar to the ART reconstruction after 100 iterations. The latter is an approximation to the pseudoinverse solution and so, in the limit of small σ_n , the MAP approaches the maximum-likelihood estimate. When σ_n is somewhat larger than the actual rms noise, the MAP solution becomes a blurry, low-contrast version of the actual scene. A Fourier spectral analysis of the reconstruction with the largest assumed rms noise value indicates that, at intermediate spatial frequencies, the amount of blurring is consistent with the expected blur filter $\approx \omega^{-1}$, where ω is the radial spatial frequency.

Figure 3 shows the average amplitude of the 100 low-contrast discs and the approximate 300 disc regions taken from the background, calculated as the average within an appropriately placed 8-pixel-diameter circle. These amplitudes approach an approximation to their actual values as σ_n decreases to zero. The values for the ART reconstructions are displayed in the graph at $\sigma_n = 0$. The most accurate value for the amplitude of the discs is about 12% lower than its actual value because of the small blurring effects in the reconstruction that are not properly accounted for in using the average within a sharp-edged circle. As σ_n increases from zero, amplitude estimates of the discs and of the background move away from their actual values toward $\bar{f} = 0.0423$, which is approached asymptotically as σ_n gets large.

In Fig. 4 we see that the rms residual, proportional to the square root of χ^2 , increases as σ_n increases, slowly at first, and then more rapidly. The minimum value is 1.70, only slightly lower than the rms noise value (known to be 1.94) indicative of a sufficient number of measurements to perform the reconstruction. This figure also displays the detectability index d' obtained at various σ_n values. The detectability index for the low-contrast discs remains essentially constant until the assumed rms noise exceeds its actual value, but eventually decreases for large σ_n . This decrease probably arises because of the expected and observed

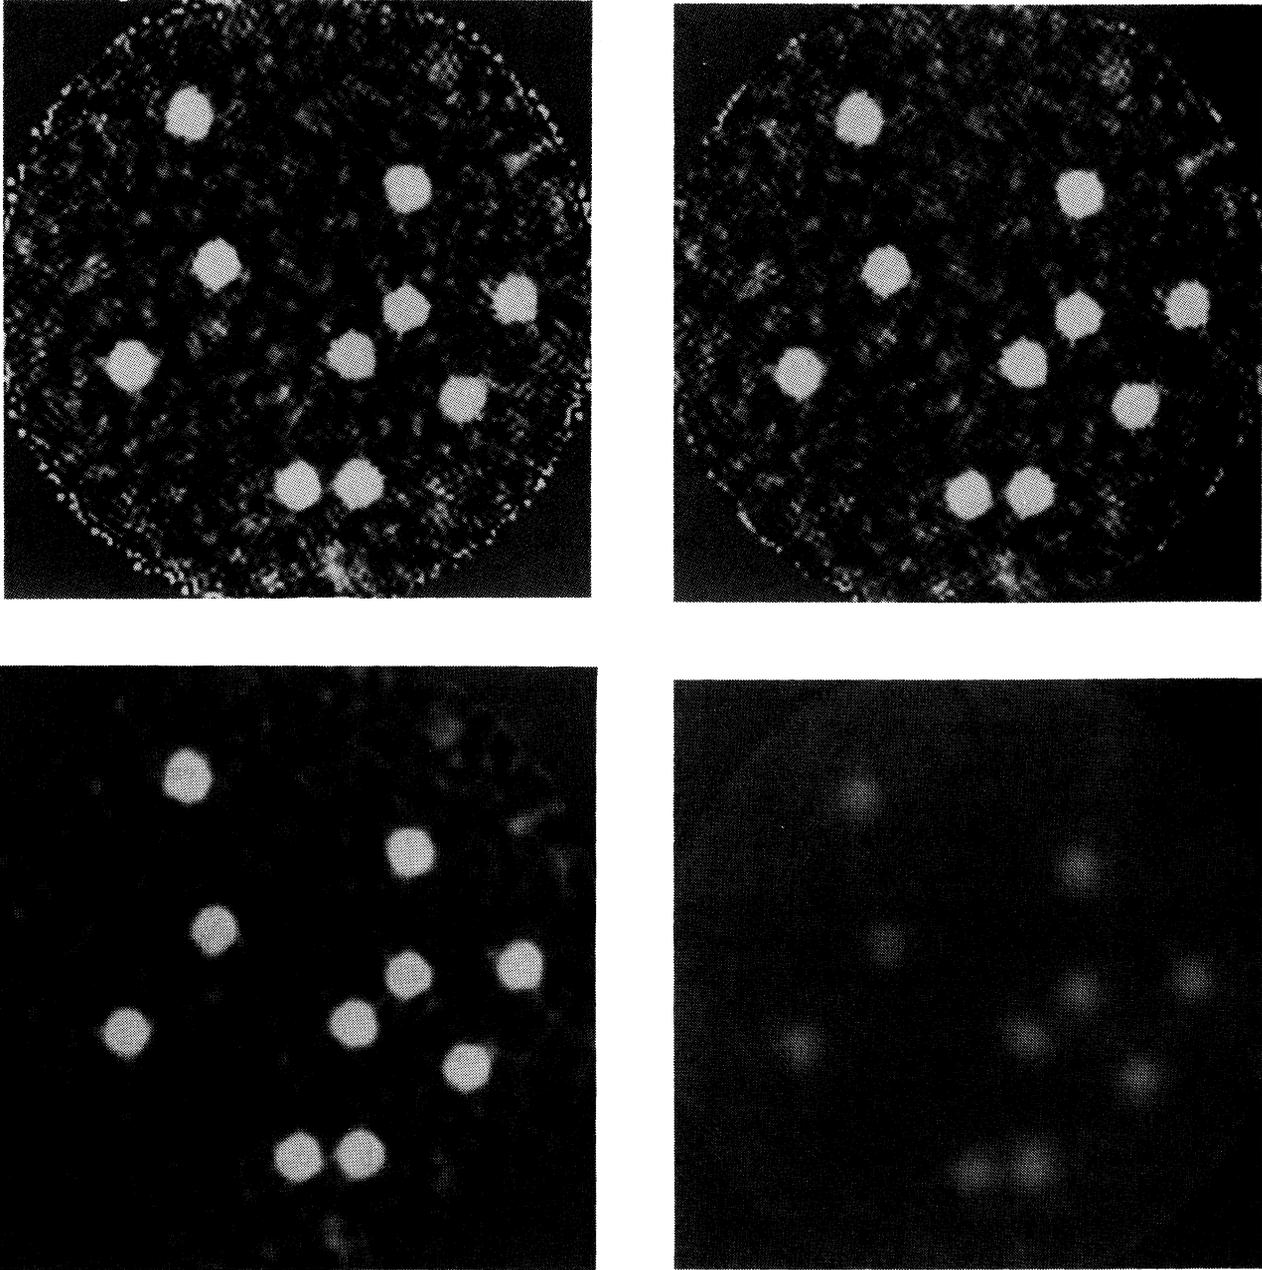


Figure 2: Linear MAP reconstructions from the same data for different assumed rms noise values σ_n of 0.1 (upper right), 2.5 (lower left), and 10 (lower left). The projection data consist of 100 views with added noise with an actual (fixed) rms value of 1.9. For comparison, the unconstrained ART reconstruction (100 iterations) obtained from the same data (upper left). All images are displayed with the same window of -0.2 to 0.4.

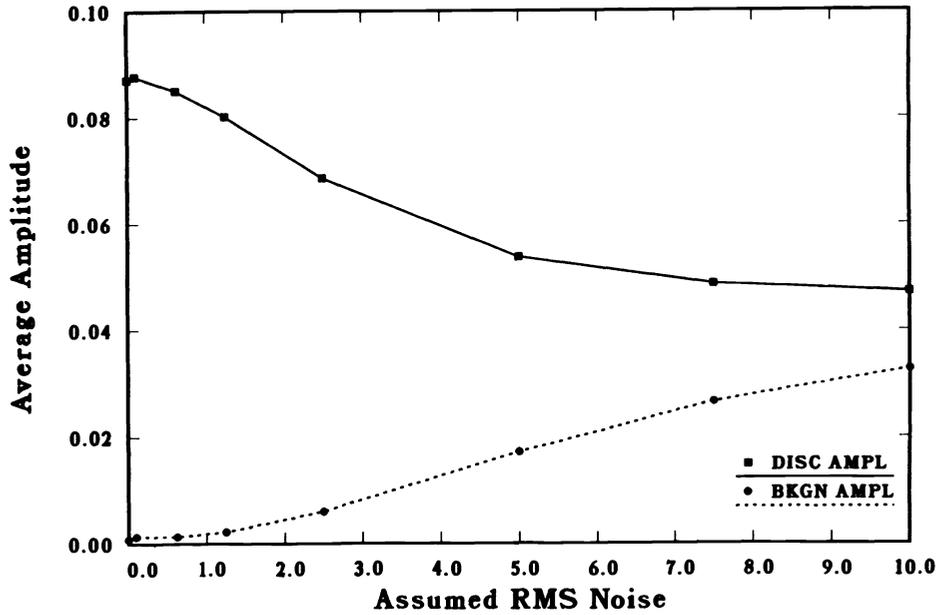


Figure 3: The average amplitude measured for the discs and background region in MAP reconstructions obtained with various assumed rms noise values for 100 views actually contaminated with an rms noise value of 1.9. For this and subsequent graphs, the leftmost point of each curve is for the ART reconstructions.

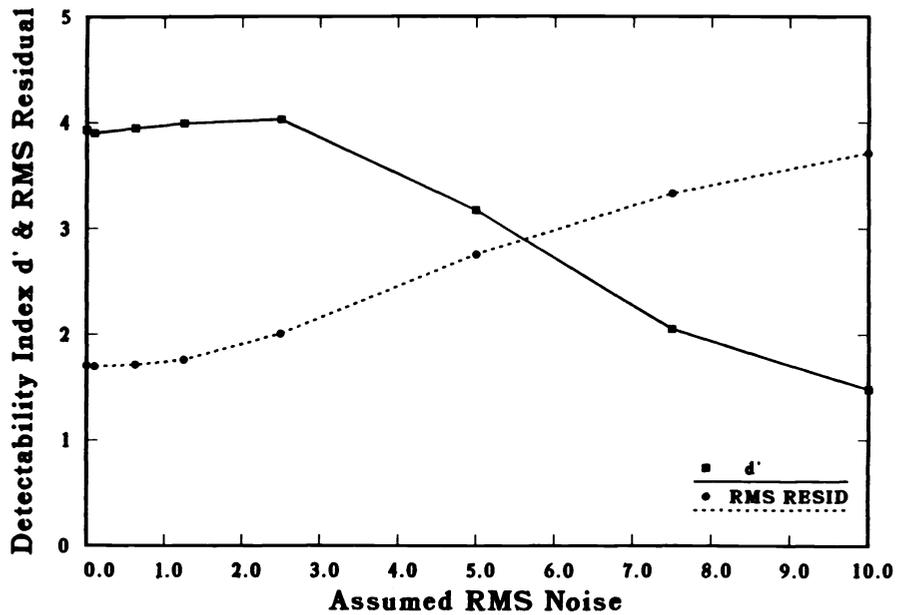


Figure 4: The detectability index and the rms residual obtained in MAP reconstructions from 100 views with an rms noise value of 1.9.

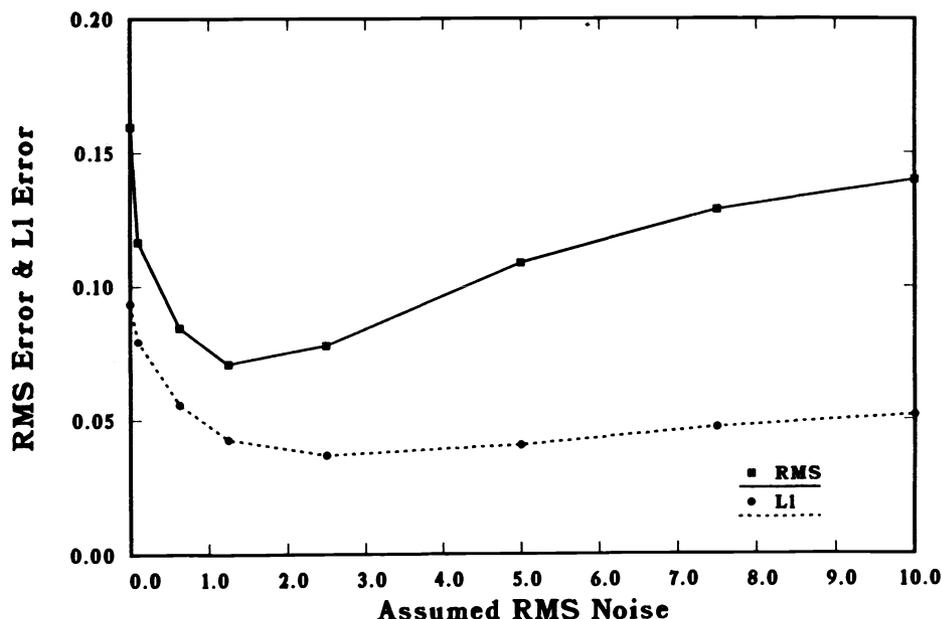


Figure 5: The root mean square (L2) and L1 error in the reconstructions based on the difference between the reconstruction and the original scene, for MAP reconstructions from 100 views with an rms noise value of 1.9.

characteristic blurring of the linear MAP algorithm for large σ_n . For the performance of the detection task, it does not matter what value is assumed for the rms noise, so long as it is less than or equal its actual value. We note that the detectability index based on the area under the receiver operating characteristic (ROC) curve d_A [1] is always consistent with d' for the examples shown here. This is not surprising because we expect the underlying histograms in the decision variable to be Gaussian shaped for these unconstrained reconstructions.

Figure 5 shows both the L2 and L1 norm errors in the MAP reconstructions as a function of the assumed rms noise value. A fairly well-defined minimum exists in the rms (L2 norm) error at an assumed noise value that is slightly below its actual value. As σ_n approaches zero, both errors increase dramatically because of the presence of high-frequency noise in the corresponding reconstructions. What is interesting is that the performance of the specified detection task is not adversely affected by this high-frequency noise. One of these two measures of reconstruction error is often used to pick the appropriate σ_n that 'optimizes' the reconstruction algorithm. We observe that, at least as far as detection of medium-sized objects is concerned, there is no corresponding deterioration in task performance for small σ_n . It remains to be seen whether other tasks dependent on the information content at higher spatial frequencies, such as the Rayleigh-inspired task of binary/singlet discrimination [21], are adversely affected by such noise.

Figures 6 and 7 show that results of similar character are obtained for a limited angle situation, 16 views covering 180° with only a small amount of noise added (rms value = 0.5). The MAP reconstructions are obtained with 100 iterations, which provided good convergence for all but the smallest σ_n value (of 0.1). The ART reconstruction, whose results are presented on the graphs at $\sigma_n = 0$, is obtained with 150 iterations and the relaxation parameters $\lambda_0 = 1.5$, $r_\lambda = 0.980$. The ART reconstruction is virtually identical in appearance to the latter MAP reconstruction. A significant change from the overdetermined reconstructions is the approximate 25% drop in the estimated amplitude of the discs and a decrease in the background amplitude in the small σ_n limit. This effect is caused by the lack of determinacy of the solution arising from the insufficient number of available views, which the algorithm counteracts by pushing the solution in the direction of the ensemble mean $\bar{f} = 0.0423$. The convergence toward this constant for

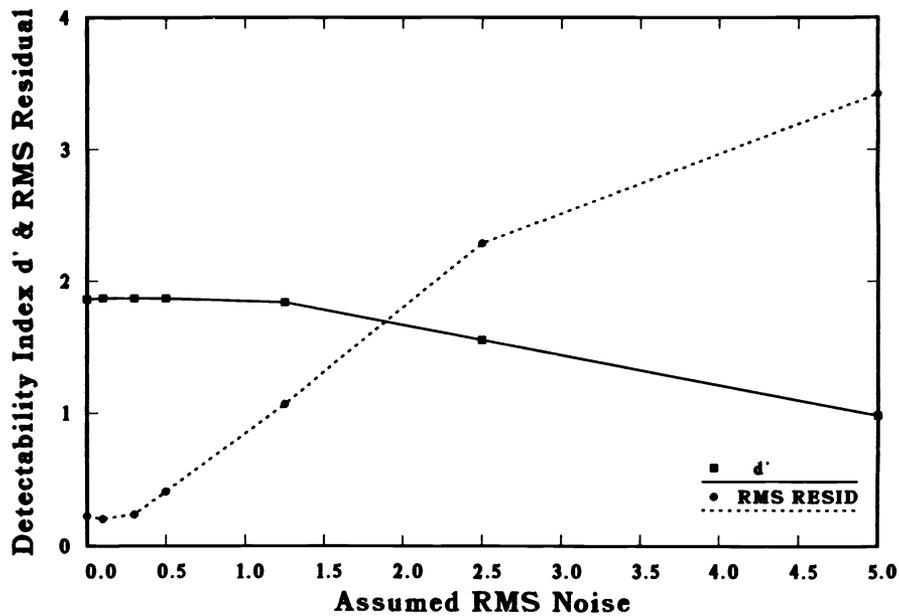


Figure 6: The average amplitude measured for the discs and for the background region in MAP reconstructions obtained with various assumed rms noise values for 16 views contaminated with an actual rms noise value of 0.5. The leftmost point of each curve is for the ART reconstructions.

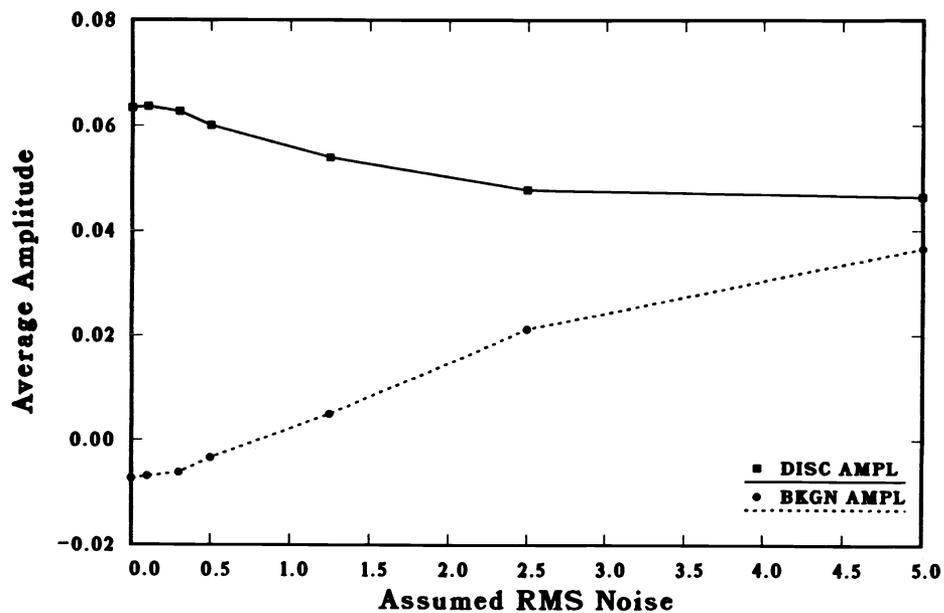


Figure 7: The detectability index for the low-contrast discs and the rms residual obtained in MAP reconstructions from 16 views with an rms noise value of 0.5.

large σ_n is similar to that observed above for 100 views.

Another departure from the earlier results for 100 views is that the minimum rms residual is 0.20, significantly smaller than the known rms noise value of 0.5. This dramatic reduction is again caused by the limited number of views. In this underdetermined problem, it should be possible to drive the residuals to zero. The fact that this lower limit is not reached indicates the existence of a small amount of regularization in the solution. Without such regularization, the solutions could be very noisy indeed. The detectability index behaves as before.

Discussion

We have investigated the performance of two related tasks: object detection and object amplitude estimation. These tasks are related because the best amplitude estimate is the appropriate decision variable for the detection task. However, we observe different results for these two tasks as a function of σ_n . It is found that in MAP reconstructions of scenes containing a mixture of high- and low-contrast discs, the contrast of the low-contrast discs relative to the background decreases steadily as σ_n increases. As $\sigma_n \rightarrow \infty$, the solutions approach the constant used for the ensemble average \bar{f} . The estimates for the amplitude of these discs can deviate substantially from their actual values. On the other hand, the detectability index does not change as quickly. The reason is that detectability is based on the difference between the estimated amplitude of the object and the estimated background value compared to their rms deviations. The choice of σ_n , which is akin to stopping an iterative algorithm at a predetermined value of χ^2 , affects the declination in the estimated amplitudes from their true values.

The similarity between the equations on which Tikhonov regularization (7) and linear MAP (1) are based, suggests that the same type of behavior noted above for MAP will be exhibited by variations of the Tikhonov approach. Some such variants include the standard Tikhonov form in which the regularization term is simply $|f|^2$ (i. e., minimum norm), and the maximum entropy (MAXENT) algorithm [8] in which the regularization term is $\sum_i f_i \ln(f_i/e\bar{f}_i)$, where \bar{f} is the image that MAXENT will produce in the limit as the rms noise in the data goes to ∞ , aptly called the default image. Early work on the evaluation of the MAXENT algorithm indicates that it behaves similarly to the linear MAP method in some respects [22].

Consideration of how the linear MAP algorithm functions, as seen from Eq. (3), leads to the observation that, on a pixel-by-pixel basis, unless $\bar{f}_i = f_i$, for nonzero σ_n the linear MAP estimate will be biased away from the actual value of the original image f_i towards \bar{f}_i . This tendency increases as σ_n increases. It is generally observed in the field of regularization that there is a tradeoff between spatial resolution, which controls the fluctuations in the reconstruction and the conditional bias. This behavior seems to [14,15] be unavoidable and we suggest it is an inevitable property of solutions to ill-posed problems. The bias effect should be kept in mind by those who wish to derive quantitative results from image reconstructions. The desire to eliminate bias tends to support the choice of a relatively small value of σ_n . An interesting aspect of the present work is that neither the average amplitude nor the detectability is affected very much by choosing very small values for σ_n . Conventionally the choice of σ_n is made to minimize the rms reconstruction error, which would advocate σ_n be close to σ_0 , the actual rms noise. Minimum L1 error would speak for even higher σ_n . As seen in Fig. 3, a significant decrease (about 20%) in the contrast between the disc amplitudes and the background accompanies these conventional choices. This potential deviation from reality should be borne in mind by those who make use of algorithms that employ a constraint on χ^2 . The same arguments apply to the other mechanisms used to regularize solutions to ill-posed problems mentioned in the introduction.

Acknowledgements

I express my appreciation to the following people for their continued interest in this work and for their worthwhile suggestions: Harry H. Barrett, Stephen F. Gull, Kyle J. Myers, Didier M. Saint-Felix,

and Robert F. Wagner. I have also received enlightening guidance from Walter Carrington, Vance Faber, and Tom Manteuffel in conjugate directions.

References

- [1] K. M. Hanson. Method to evaluate image-recovery algorithms based on task performance. *Proc. SPIE*, 914:336–343, 1988.
- [2] K. M. Hanson. Method to evaluate image-recovery algorithms based on task performance. Submitted to *J. of the Optical Soc. of Amer. A*, 1990.
- [3] A. P. Sage and J. L. Melsa. *Estimation Theory with Applications to Communications and Control*. Robert E. Krieger, Huntington, 1979.
- [4] E. Veklerov and J. Llacer. Stopping rule for the MLE algorithm based on statistical hypothesis testing. *IEEE Trans. Med. Imaging*, MI-6:313–319, 1987.
- [5] R. Gordon, R. Bender, and G. Herman. Algebraic reconstruction techniques for three-dimensional electron microscopy and x-ray photography. *J. Theor. Biol.*, 29:471–481, 1970.
- [6] P. Gilbert. Iterative methods for the three-dimensional reconstruction of an object from projections. *J. Theor. Biol.*, 36:105–117, 1972.
- [7] H. C. Andrews and B. R. Hunt. *Digital Image Restoration*. Prentice-Hall, Englewood Cliffs, New Jersey, 1977.
- [8] S. F. Gull and G. J. Daniell. Image reconstruction from incomplete and noisy data. *Nature*, 272:686–690, 1978.
- [9] R. R. Bevington. *Data Reduction and Error Analysis for the Physical Sciences*. McGraw-Hill, New York, 1969.
- [10] A. D. Whalen. *Detection of Signals*. Academic, New York, 1971.
- [11] K. M. Hanson. Bayesian and related methods in image reconstruction from incomplete data. In Henry Stark, editor, *Image Recovery: Theory and Application*, pages 79–125, Academic, Orlando, 1987.
- [12] A. Albert. *Regression and the Moore-Penrose Pseudoinverse*. Academic, New York, 1972.
- [13] A. N. Tikhonov and V. Y. Arsenin. *Solutions of Ill-Posed Problems*. Winston-Wiley, New York, 1977.
- [14] M. Z. Nashed. Operator-theoretic and computational approaches to ill-posed problems with applications to antenna theory. *IEEE Trans. Antennas Propagat.*, AP-29:220–231, 1981.
- [15] D. M. Saint-Felix. *Restauration d'image: Régularisation d'un problème mal-posé et algorithmique associée*. PhD thesis, Laboratoire des Signaux et Systèmes, École Supérieure d'Électricité, Gif-Sur-Yvette, France, 1987.
- [16] K. M. Hanson and G. W. Wecksung. Bayesian approach to limited-angle reconstruction in computed tomography. *J. Opt. Soc. Amer.*, 73:1501–1509, 1983.
- [17] P. E. Gill, W. Murray, and F. H. Wright. *Practical Optimization*. Academic, New York, 1981.
- [18] A. V. Oppenheim, A. S. Willsky, and I. T. Young. *Signals and Systems*. Prentice-Hall, Englewood Cliffs, New Jersey, 1983.
- [19] K. M. Hanson. Optimization for object localization of the constrained algebraic reconstruction technique. *Proc. SPIE*, 1090:146–153, 1989.
- [20] K. M. Hanson. POPART - Performance OPTimized Algebraic Reconstruction Technique. *Proc. SPIE*, 1001:318–325, 1988.
- [21] K. M. Hanson. Optimization of the constrained algebraic reconstruction technique for a variety of visual tasks. *Proc. Information Processing in Medical Imaging*, Berkeley, CA, June 19-23, 1989.
- [22] K. J. Myers and K. M. Hanson. Comparison of the algebraic reconstruction technique with the maximum entropy reconstruction technique for a variety of detection tasks. To be published in *Proc. SPIE*, 1231:, 1990.